

**METHOD, SYSTEM AND PROGRAM PRODUCT FOR ACTIVELY  
MANAGING CENTRAL QUEUE BUFFER ALLOCATION USING A  
BACKPRESSURE MECHANISM**

**Cross-Reference to Related Applications**

**[0001]** This application contains subject matter which is related to the subject matter of U.S. Patent Application Serial No. \_\_\_\_\_, by Garmire, et al., entitled “METHOD, SYSTEM AND PROGRAM PRODUCT FOR ACTIVELY MANAGING CENTRAL QUEUE BUFFER ALLOCATION” (Attorney Docket No. POU920030161US1), which is assigned to the same assignee as this application and filed on the same day as this application, and which is hereby incorporated herein by reference in its entirety.

**Technical Field**

**[0002]** This invention relates in general to packet flow control through a data transfer network switching node, and more particularly, to techniques for facilitating packet flow by actively managing allocation of space in a node’s central queue allotted to switching node ports based on the amount of unused space currently available in the central queue.

**Background of the Invention**

**[0003]** A conventional switching node for a data communications (or transfer) network utilizes dedicated buffers at its input ports or at its output ports. These dedicated buffers couple the ports to the switching logic that directs data packets received by the input ports to the appropriate output ports and facilitate flow control of the data packets transmitted through the data communications network. Common data traffic patterns in a data transfer network often result in widely disparate volumes of traffic at the ports of network switching nodes, and these disparate traffic volumes typically vary over time. Consequently, the use of buffers which are dedicated to respective ports in the network

switching nodes is often not well-matched to the data traffic characteristics. Input port blocking can occur in a busy port because its buffer fills to capacity while the data packet at the front of the buffer's queue waits to be switched to a busy output port. Meanwhile, the buffers of other ports that are only receiving a small amount of traffic have free space available that cannot be utilized.

### Summary of the Invention

**[0004]** The shortcomings of the prior art are overcome and additional advantages are provided through the provision of a method of packet flow control for a switching node of a data transfer network. The method includes actively managing space allocations in a central queue of a switching node allotted to the ports of the switching node. The active management of the space allocations in the central queue is based on the amount of unused space currently available in the central queue and an amount of currently-vacant storage space in a storage device of a port.

**[0005]** A further aspect of the technique of managing space allocations in the central queue includes separately tracking unused space that is currently unallocated, and unused space that is allocated to ports of the switching node but has been vacated due to a packet being transferred from the central queue to an output port. Each port is offered vacated space that is currently allocated to it and a quantity of the currently unallocated space in the central queue (e.g., approximately periodically). Yet another aspect of the active management of space allocations in the central queue to the ports comprises dynamically allocating the quantity of unused space offered to a port to one or more virtual lanes of the port.

**[0006]** Systems and computer program products corresponding to the above-summarized methods are also described and claimed herein.

[0007] Additional features and advantages are realized through the techniques of the present invention. Other embodiments and aspects of the invention are described in detail herein and are considered a part of the claimed invention.

#### Brief Description of the Drawings

[0008] The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the claims at the conclusion of the specification. The foregoing and other objects, features, and advantages of the invention are apparent from the following detailed description taken in conjunction with the accompanying drawings in which:

[0009] FIG. 1 illustrates one embodiment of a switching node of a data transfer network employing packet flow control, in accordance with an aspect of the present invention;

[0010] FIG. 2 is a flowchart of one aspect of a technique for managing central queue buffer allocations to the ports of a switching node of a data transfer network, in accordance with an aspect of the present invention;

[0011] FIGS. 3a and 3b illustrate a flowchart of another aspect of a technique for managing central queue buffer allocations to the ports of a switching node, in accordance with an aspect of the present invention, wherein space in the central queue assigned to a port is further allocated to virtual lanes of the port;

[0012] FIG. 4 illustrates one embodiment of port control logic for the switching node embodiment of FIG. 1, in accordance with an aspect of the present invention;

[0013] FIG. 5 illustrates one embodiment of virtual lane calculators for the port control logic embodiment of FIG. 4, in accordance with an aspect of the present invention; and

[0014] FIG. 6 illustrates a flowchart of one embodiment of the arithmetic logic of FIG. 5, in accordance with an aspect of the present invention.

**Best Mode for Carrying Out the Invention**

[0015] Generally stated, described herein is a packet flow control technique for a switching node of a data transfer network. In one aspect, the technique actively and dynamically manages space allocations in a central queue of a switching node that are given to the ports of the switching node to be used for packet buffering. The active and dynamic management of the space allocations in the central queue is based on the amount of currently unused space available in the central queue. In another aspect, the technique includes allotting the space allocation given to a port among one or more virtual lanes of the port. Space in the central queue is allocated in n-byte blocks (a.k.a. "chunks"). Each block is represented by one credit. Based on the amount of central queue space that is allotted to a port's virtual lanes, the port can grant a corresponding number of credits to the transmitting end of the link for flow control purposes.

[0016] One embodiment of a switching node of a data transfer network employing packet flow control, in accordance with an aspect of the present invention, is described below with reference to FIG. 1. This embodiment of a switching node comprises multiple input ports (110, 120, and 130) coupled to a central queue 140 and multiple output ports (160, 170, and 180) also coupled to central queue 140. The switching node embodiment illustrated in FIG. 1 also includes central queue control logic 150. Central queue control logic 150 is coupled to central queue 140, input ports 110, 120, and 130, and output ports 160, 170, and 180.

[0017] Input port 110 interfaces to a communications link via a receiver and link protocol engine unit 112. Receiver and link protocol engine unit 112 is coupled to a first-in, first-out (FIFO) register 114 and port control logic 116. FIFO register 114 of input port 110 is coupled to central queue 140. Input ports 120 and 130 are similarly configured.

[0018] FIFO register 162 of output port 160 interfaces to central queue 140 and transmitter and link protocol engine unit 164 of output port 160. Transmitter and link protocol engine unit 164, in turn, is coupled to port control logic 166 and a communications link. Output ports 170 and 180 are similarly configured.

[0019] The following discussion summarizes the flow of a data packet through switching node embodiment 100 illustrated in FIG. 1 and the interaction of various of its constituent parts. By way of example, the flow of a data packet through switching node 100 where the data packet is received by input port 110 and is transmitted eventually via output port 160 is assumed. After the data packet is received by receiver and link protocol engine unit 112, it is placed in FIFO register 114. From FIFO register 114, the packet is buffered in central queue 140. Port control logic 116 decrements the space available to the virtual lane to which the packet is assigned by an amount equal to the size of the packet. A port determines a packet's virtual lane from an attribute of the packet (e.g. the packet's service level field) and a service-level-to-virtual-lane mapping table.

[0020] When the destination output port is not busy transmitting another packet, the packet buffered in central queue 140, can be transferred to FIFO register 162 of output port 160. The data packet is read from FIFO register 162 for transmission over a link by transmitter and link protocol engine unit 164. The transfer of the received packet from central queue 140 to FIFO register 162 in preparation for transmission causes space which is allocated to a virtual lane of a port to be vacated. This vacated space is offered back to the virtual lane of the port from which the packet was received by central queue control logic 150 in order to maintain an allocation of space in central queue 140 for all of the multiple ports.

[0021] In addition, the space allocations in central queue 140 for the ports of the switching node are actively managed based on the amount of unused space available in central queue 140. The unused space in central queue 140 includes space that is not currently allocated to any of the switching node's ports and vacated allocated space, i.e. space that is currently allocated to a virtual lane of a port and had been used to buffer a

received packet, but has been vacated as the result of the packet being transferred to a FIFO register of an output port. Vacated allocated space can also result from an input port discarding a received packet.

**[0022]** In one embodiment of a technique for actively managing space allocations in central queue 140 for the ports of the switching node, the management is distributed among central queue control logic 150 and the port control logic of multiple ports. FIG. 2 is a flowchart 200 of one embodiment of central queue control logic for allocating space in the central queue to the ports of the switching node embodiment shown in FIG. 1, in accordance with an aspect of the present invention. With reference to FIG. 2, several parameters utilized in the management of space allocations are initialized 210 to begin the process. These include: SharedCredits, which is the amount of unallocated space available in the central queue measured in units of credits, OfferRC, which is the amount of vacated allocated space measured in units of credits, and MaxOffer, which is the maximum amount of unallocated space available in the central queue that central queue control logic 150 of FIG. 1 can offer to a port per management interval.

**[0023]** Following initialization of parameters, active central queue space allocation management begins with selecting a port to which to offer credits, i.e. an amount of unused space, using a defined rule 212. One example of such a defined rule is selecting each port in turn approximately periodically so that central queue space is fairly allocated among the ports of the switching node. The central queue control logic separately tracks the amount of vacated allocated space for the virtual lanes of each of the ports and the amount of unallocated space currently available in the central queue.

**[0024]** As indicated in the flowchart of FIG. 2, these two categories of unused space in the central queue are allocated differently. The counter table OfferRC comprises a counter for each virtual lane. The counter for each virtual lane of the selected port is set equal to an amount of space in the central queue corresponding to all vacated allocated space assigned to the associated virtual lane 220. The counter variable SharedCredits is used to track the amount of unallocated space currently available in the central queue. If

a test of the counter variable SharedCredits indicates that there is unallocated space in the central queue 230, an amount of unallocated space in the central queue, up to an amount defined by the parameter MaxOffer, is allocated to the selected port by setting the variable OfferSC equal to the smaller of the following two amounts of space: MaxOffer and SharedCredits 232, 233, and 234. The amount of space allocated to the selected port is also deducted from the amount of unallocated space currently available in the central queue 233 and 234. The amount of unused space represented by the values of the counter table OfferRC and variable Offer SC is offered to the selected port 240. The number of credits returned by each port, representing an amount of previously allocated central queue space, is added to the pool of unallocated space in the central queue 250 before the space management process of the central queue control logic returns to the beginning of the process loop 212.

**[0025]** As discussed above, two categories of unused space in the central queue – vacated allocated space and unallocated space – are tracked separately and allocated differently. The reasons for this are as follows. On one hand, vacated allocated space is offered by the central queue control logic to the virtual lane of the port to which it is currently allocated to ensure that each virtual lane receives a designated minimum amount of buffer space, thereby preventing those virtual lanes from suffering blocking due to heavy traffic and blocking problems experienced in other virtual lanes. On the other hand, offered shared credits are offered on a fair basis to all ports in order to dynamically provide additional buffer space in the shared central queue to various ports and their respective virtual lanes as required to serve time-varying traffic to the ports of the switching node.

**[0026]** Further aspects of a technique for actively managing space allocations in the central queue of a switching node for the node's ports, in accordance with an aspect of the present invention, are described below with reference to the flowchart shown in FIGS. 3a and 3b. This flowchart illustrates one embodiment of a technique for allocating space in the central queue (which has been offered to a port), to the virtual lanes of the

input port of the switching node embodiment shown in FIG. 1. With reference to FIG. 3a, several variables utilized in the allocation of space offered to a port are initialized 310 to begin the process. These include the following variables for each virtual lane: CurrentCredit(VL), which counts the amount of space currently available (unused) in the allocation to that virtual lane; DesiredCredit(VL), which represents the target amount of unused buffer space in the central queue that the port control logic will attempt to maintain for the virtual lane, working in coordination with the central queue control logic; and DebtCredit(VL), which represents the amount of buffer space in the central queue that has been borrowed by the virtual lane from the pool of unallocated space available in the central queue (but not yet returned). Initially, DebtCredit(VL) is set equal to -DesiredCredit(VL). The negative value of DebtCredit(VL) is an indication that this virtual lane is owed space in the central queue.

**[0027]** After initialization, the repeating loop in the space allocation technique for the virtual lanes begins with selecting the virtual lane (VL) of the port that is currently in the greatest need of buffer space in the central queue 312. In one example, the VL most in need of space in the central queue is determined to be the virtual lane having the largest difference between DesiredCredit(VL) and CurrentCredit(VL). In the flowchart shown in FIGS. 3a and 3b, buffer space in the central queue is measured and tracked in units called “credits”.

**[0028]** Once a virtual lane has been selected 312, the remaining steps in the repeating loop of the virtual-lane space allocation technique involve tracking various measures of space allocation and usage as well as space allocation decisions that are directed to the selected VL. First, DebtCredit(VL) is tested to determine whether the selected virtual lane is owed any space in the central queue 320. In one embodiment, a virtual lane is owed credits when the value of DebtCredit is negative. If the selected VL is owed any credits, OfferRC(VL) is compared to (DesiredCredit(VL) – CurrentCredit(VL)) 321 where (DesiredCredit(VL) – CurrentCredit(VL)) is the difference between the target amount of unused buffer space in the central queue for virtual lane VL and the current

amount of unused space in the allocation to that virtual lane. If OfferRC(VL) is less than or equal to ((DesiredCredit(VL) – CurrentCredit(VL)), all OfferRC credits offered to the port for the selected VL are allocated to that virtual lane 322. Alternatively, if OfferRC(VL) is greater than ((DesiredCredit(VL) – CurrentCredit(VL)), the port control logic keeps only enough of the OfferRC(VL) credits for the selected virtual lane to increase the current amount of unused space in the selected virtual lane's central queue allocation (CurrentCredit(VL)) to the target amount ((DesiredCredit(VL)) 323. In this case, the remaining OfferRC(VL) credits, i.e. OfferRC(VL) – ( DesiredCredit(VL) – CurrentCredit(VL) ), are returned to the pool of unallocated space in the central queue by setting the variable CQCreditReturn equal to the number of remaining OfferRC(VL) credits 323. As discussed previously with reference to FIG. 2, OfferRC credits offered to a port for a particular virtual lane by central queue control logic correspond to the amount of vacated allocated space in the central queue that is currently assigned to that virtual lane of the port.

**[0029]** If the selected virtual lane is not owed any credits, the allocation processing of port control logic goes on to determine how much of the vacated allocated space in the central queue offered to the port for the selected virtual lane (i.e. the OfferRC credits for the selected virtual lane), to allocate to the selected VL and how the determined allocation to the selected VL is utilized. In order to simplify the foregoing discussion, vacated allocated space in the central queue offered to a port for a virtual lane will be referred to as rebated space or rebated credits, and OfferRC(VL) will designate the counter in table OfferRC for the selected virtual lane. If the rebated space exceeds the current central-queue space debt of the selected virtual lane 330, the rebated credits are first used to eliminate the VL's debt by returning the number of credits, corresponding to the current balance of central-queue space borrowed, to the pool of unallocated space in the central queue. In one example, the variable CQ\_CreditReturn is set to DebtCredit(VL) to effectuate satisfaction of the current space debt 332. Also, DebtCredit(VL) set equal to zero 332 to indicate that the virtual lane no longer has a debt.

**[0030]** After satisfying the selected VL's space debt, the logic determines whether the number of remaining rebated credits, i.e. those not used to satisfy the space debt, exceed the number required to raise the amount of currently available (unused) space in the VL's allocation to the desired amount 334. If so, the number of remaining rebated credits required to increase the amount of space available in the VL's allocation in the central queue to the desired amount is allocated to the selected virtual lane 336. The number of remaining rebated credits exceeding the amount required is returned to the pool of unallocated space in the central queue by adding the excess credits to CQ\_CreditReturn 336. If test 334 is negative, all offered rebated credits are allocated to the selected virtual lane 335. In either case, the count kept in variable CurrentCredit(VL) is incremented as appropriate to reflect the increase in currently available space in the VL's allocation (334 and 335).

**[0031]** If, however, the rebated space does not exceed the current central-queue space debt of the selected virtual lane 330, all of the rebated credits offered to the port are returned to the pool of unallocated space in the central queue in order to reduce the selected virtual lane's space debt 331. In one example, this space debt reduction is effectuated by setting the variable CQ\_CreditReturn to OfferRC(VL) and deducting OfferRC(VL) credits from the space debt counter variable DebtCredit(VL).

**[0032]** After allocating any rebated space that has been offered to the port for the selected virtual lane, the management of space allocations proceeds with allocating the quantity of space from the pool of unallocated space in the central queue that has been offered to the port. This quantity of offered space is indicated by the value of the variable OfferSC. To simplify the following discussion, the quantity of space from the pool of unallocated space in the central queue that has been offered to the port will also be referred to as offered shared credits. The allocation of offered shared credits among the port's virtual lanes begins with determining whether any shared credits have been offered to the port 340. If not, processing returns to the beginning of the loop to select the virtual lane in the greatest need of credits currently (312, FIG. 3a). If a quantity of space from

the pool of unallocated space in the central queue that has been offered to the port, OfferSC is tested to determine whether the offered shared credits exceed the number of credits required to increase the selected virtual lane's available space in its central queue allocation to its target level 342. If it does not, all offered shared credits are allocated to the selected virtual lane's space allocation in the central queue 344; the selected VL's debt is increased by the number of offered shared credits 344; and processing proceeds to the selection of the virtual lane currently having the greatest need for space in the central queue 312 in FIG. 3a.

[0033] Alternatively, if the offered shared credits exceed the number of credits required to increase the selected virtual lane's available space in its central queue allocation to its target level, the number of offered shared credits required to raise the virtual lane's available space in the central queue to the target amount is allocated to the selected virtual lane's space allocation in the central queue 346; the selected VL's debt is increased by the number of offered shared credits so allocated to the VL 346; the excess offered shared credits beyond what is needed to raise the VL's available credits in the central queue to the target amount are returned to the pool of unallocated space available in the central queue 346; and processing proceeds to the selection of the virtual lane currently having the greatest need for space in the central queue 312 in FIG. 3a.

[0034] It should be noted that because DesiredCredit(VL) is a factor in the test to determine which virtual lane has the greatest need for credits 312 and also influences how many of the offered shared credits are allocated to the selected virtual lane (342, 344, and 346) for each pass through the repeating loop in the technique for allocating space in the central queue offered to a port of a switching node, it is possible to dynamically change the space allocations in the central queue for the virtual lanes of a port without having to reinitialize the switching node by simply changing the value of the DesiredCredit(VL) parameter.

[0035] One embodiment of port control logic 116 (also referred to as port credit manager) for the switching node embodiment of FIG. 1, in accordance with an aspect of

the present invention, is described below with reference to FIG. 4. By way of example, the embodiment of FIG. 4 partitions the port's resources into four virtual lanes, which are designated by indices 0, 1, 2, and 3. Port control logic 116 receives several signals that it uses to manage a port's space allocation in central queue 140 of FIG. 1 and allocate space assigned to the port among multiple virtual lanes. These input signals include DesiredCredit(0:1)(0:11), CQ\_Grant\_Credit, IP\_Decr\_Credit(0:3), OP\_Rebate\_Credit(0:1)(0:1), and IP\_Incr\_Credit(0:3). In most cases, the first set of indices references the VL, and the second set is the magnitude. The signal DesiredCredit designates the target quantity of available (i.e. unused) buffer space to maintain in the central queue for each virtual lane. CQ\_Grant\_Credit is a quantity of unallocated space in the central queue that is currently being offered to the port (i.e., OfferSC from FIG. 2). The group of signals IP\_Decr\_Credit(0:3) indicates an amount of space used to buffer packets received by the port for each of the virtual lanes since the previous update, which, therefore, must be deducted from the virtual lanes' available space allocations in the central queue. OP\_Rebate\_Credit indicates the amount of space in the central queue that is being offered as a rebate to the port's VL (i.e., OfferRC from FIG. 2) because one or more packets previously received by the port and buffered in the central queue were transferred from the central queue to an output port since the previous update.

IP\_Incr\_Credit indicates the amount of space in the central queue that is offered as a rebate to the port because one or more packets previously received by the input port were discarded by the port itself in accordance with the protocol employed by the data transfer network.

**[0036]** Signals CQ\_Grant\_Credit, IP\_Decr\_Credit(0:3), OP\_Rebate\_Credit(0:1)(0:1), and IP\_Incr\_Credit(0:3) are passed to virtual lane credit calculators 420, 421, 422, and 423 by virtual lane decoder logic 410. In particular, virtual lane decoder logic 410 passes the components of signals IP\_Decr\_Credit(0:3), IP\_Incr\_Credit(0:3), and OP\_Rebate\_Credit identified by indices 0, 1, 2, and 3 to the associated virtual lane credit calculators 420, 421, 422, and 423, respectively. On the other hand, signal CQ\_Grant\_Credit is directed by demultiplexer 411 to one of the VL credit calculators via

virtual lane decoder logic 410 in response to the signal Next\_VL, which is generated by arbiter logic 430. Also, the DesiredCredit signal component associated with each virtual lane is provided to the corresponding VL credit calculator, and all VL credit calculators receive the Allow\_Credit signal generated by scaling logic 450. The signal Allow\_Credit is a backpressure mechanism that is sent to all virtual lane calculators to prevent an overflow of a port's FIFO register and will be described in more detail below.

**[0037]** Virtual lane credit calculators 420, 421, 422, and 423 provide signals NeedCredit(0), NeedCredit(1), NeedCredit(2), and NeedCredit(3), respectively, to arbiter logic 430. The NeedCredit signals indicate the additional central queue space needs of the port's virtual lanes. Arbiter logic 430 will pick the virtual lane in the most need of credit (i.e., the virtual lane having largest value of NeedCredit). The virtual lane credit calculators also produce the CurrentCredit and CQ\_CreditReturn signals. CurrentCredit represents the current amount of unused space available in the central queue space that is currently allocated to a virtual lane. CQ\_CreditReturn represents the amount of space that is returned by a virtual lane to the pool of unallocated central-queue space. A virtual lane's CurrentCredit count is decremented by the amount of space indicated by IP\_Decr\_Credit by its VL credit calculator whenever a packet is received by the input port on the specified virtual lane. The values of OP\_Rebate\_Credit and IP\_Incr\_Credit for each virtual lane will be directed to their respective credit calculators to determine whether an increase in the CurrentCredit for a virtual lane is allowed, or the credits should be returned to the central queue via the Return\_Credit\_Out signal on the CQ\_CreditReturn ring (FIG. 1, 152).

**[0038]** Space returned to the unallocated pool of space in the central queue (CQ\_CreditReturn) is added to credits returned by other ports (Return\_Credit\_In) by adder logic 460. The sum (Return\_Credit\_Out) is placed on credit return ring 152 of FIG. 1. The CurrentCredit signals for each virtual lane are provided to the receiver & link protocol engine unit of the port, e.g. receiver & link protocol engine unit 112 of input port 110 in FIG. 1. Also, summation logic 440 computes the sum of all of the virtual

lanes' CurrentCredit signals for the port to produce the TotalCurrentCredit signal, which is an input to scaling logic 450. Other inputs to scaling logic 450 include FifoRoom and CSF. The FifoRoom signal indicates the amount of unused space in the port's input FIFO 114. The CSF (Credit Scaling Factor) signal introduces an adjustment to compensate for a difference in the bandwidth to the input port and the bandwidth between the input port and central queue.

[0039] Scaling Logic 450 is an enhancement to the port control logic 116 that provides a backpressure mechanism in case the input port bandwidth exceeds the bandwidth between the input port and central queue. In one example, the credit scaling factor is calculated as follows:

$$\text{CSF} = \text{BW}_{\text{RLPE}} / (\text{BW}_{\text{RLPE}} - \text{BW}_{\text{CQ}}) \text{ where}$$

$\text{BW}_{\text{RLPE}}$  is the bandwidth to the input port, and

$\text{BW}_{\text{CQ}}$  is the bandwidth from the input port to the central queue.

[0040] The scaling logic will assert the Allow\_Credit signal only when the following equation is satisfied:

$$(\text{CSF}) (\text{FifoRoom}) > \text{TotalCurrentCredit}$$

where FifoRoom is the number of unused credits in an input port's input FIFO.

When Allow\_Credit is not asserted, the input port will not accept any credits offered on the IP\_Incr\_Credit, OP\_Rebate\_Credit, or CQ\_Grant\_Credit signals in the form of CurrentCredit until the FifoRoom increases enough (i.e., the input port's input FIFO empties enough) to satisfy the above condition. If this condition is not satisfied, the credit will be used only to reduce the selected virtual lane's space debt, which is indicated by DebtCredit (VL). If DebtCredit is 0 while Allow\_Credit is also zero (i.e., not asserted), and credit is returned to the port, the input port will return the returned credits to the CQ and allow its DebtCredit to go negative, effectuating a "loan" to the central queue until the input port's input FIFO can empty out.

[0041] One embodiment of the virtual lane credit calculators shown in the port control logic embodiment of FIG. 4, in accordance with an aspect of the present invention, is described with reference to FIG. 5. In this embodiment, the virtual lane credit calculator collects credit returned through IP\_Incr\_Credit and OP\_Rebate\_Credit in collection cache 553 and CQ\_Grant\_Credit in collection cache 552. Collection caches 552 and 553 are just adders that accumulate the number of credits returned in a small interval of time, which is equal to the time it takes for the VL credit calculator to update. The sum of credits from IP\_Incr\_Credit and OP\_Rebate\_Credit for a particular VL during the update interval is stored in OfferRC by cache 553, and the sum of credits from CQ\_Grant\_Credit during the update interval is stored in cache 552. Signals OfferRC, OfferSC, Allow\_Credit, DesiredCredit, CurrentCredit, NeedCredit, and DebtCredit all feed arithmetic logic 590, which calculates the amount of change necessary in DebtCredit (dDebt), and CurrentCredit (dCurrent). Arithmetic logic 590 also determines how much surplus credit to return to the central queue (as indicated by the CQ CreditReturn signal).

[0042] The current value of DebtCredit is computed by adder logic 530 and subtraction logic 520 as follows. Adder logic 530 adds the change in the VL's space debt dDebt to the previous value of DebtCredit. The signal dDebt is provided by arithmetic logic 590. Then, this sum from adder logic 530 is provided as an input to subtraction logic 520. Subtraction logic 520 calculates the difference between the output of adder 530 and the signal dDesired, which is the output of subtraction logic 510 and represents the change in the target space level for the virtual lane over a management update interval. The value of DebtCredit calculated by subtraction logic 520 is provided as a feedback input to arithmetic logic 590 and adder logic 530.

[0043] The current value of CurrentCredit, the space currently available in the virtual lane's space allocation in the central queue, is computed as follows. Arithmetic logic 590 provides dCurrent, a signal which indicates the change in the available space in the VL's allocation in the central queue since the last management update, as an input to adder logic 560. The other input to adder logic 560 is the previous value of CurrentCredit. The

previous value of CurrentCredit is the output of subtraction logic 570, which is fed back via a register. The output of adder logic 560 is decremented by the value of the IP\_Decr\_Credit signal by subtraction logic 570 to produce the CurrentCredit signal. As discussed previously, the IP\_Decr\_Credit signal indicates the amount of the virtual lane's central queue space allocation that has been consumed to store received packets since the last management update.

**[0044]** This embodiment of the virtual lane credit calculators calculates the output signal NeededCredit by providing input signal DesiredCredit and output signal CurrentCredit as inputs to subtraction logic 580. Subtraction logic 580 computes NeededCredit by calculating the difference (DesiredCredit – CurrentCredit). The output signal NeededCredit is also provided as feedback to arithmetic logic 590 via a register.

**[0045]** In one example, arithmetic logic 590 computes the signals dDebt, dCurrent, and CQ CreditReturn in accordance with the technique for allocating space to the virtual lanes of a port described previously with reference to the flowchart illustrated in FIGS. 3a and 3b as shown in more detail in FIG. 6. The calculations described in FIG. 6 are executed in hardware where all four virtual lane credit values are manipulated in parallel.

**[0046]** Returning to FIG. 5, the signal dDesired represents a change in the amount of DesiredCredit for a virtual lane. If DesiredCredit is increased (allocating more space in the central queue to a VL of an input port), an immediate reduction in DebtCredit is effectuated by subtractor 520. This enables the VL to accept more reserved space from the central queue. If DesiredCredit is decreased (reducing a VL's space allocation in the central queue) an immediate increase in DebtCredit is effectuated by subtractor 520 and the virtual lane calculator will begin to return credit on CQ CreditReturn to attempt to drive the amount of space used by the VL to its current allocation of space.

**[0047]** The LinkState signal supplied to cache 553 is asserted when DesiredCredit is set to 0 and the virtual lane or port is disabled. It will cause an amount of credit equal to the value of the CurrentCredit signal to be asserted as both OfferRC and IP\_Decr\_Credit,

causing the credit calculator to convert CurrentCredit into CQ CreditReturn credit and effectively returning the reserved credit of this virtual lane of the port to the central queue.

**[0048]** The desired amount of buffer space in the central queue for each virtual lane can be set differently based on the number of VLs being used by a port, the maximum transmission unit (MTU) size for the port, and the level of performance required. The packet flow control technique in accordance with the present invention allows the desired amount of credit to be changed on the fly while the port is in operation. Thus, the amount of buffer space that the input port tries to maintain for its virtual lanes can be raised or lowered as traffic conditions change or as ports are activated deactivated without having to reset the switch.

**[0049]** FIG. 6 illustrates a flowchart of one embodiment of arithmetic logic 590 of FIG. 5, in accordance with an aspect of the present invention. Values are assigned to the variables dDebt, dCurrent, and CQCreditReturn 660 in accordance with the path taken through the decision logic comprising arithmetic logic 590. The processing begins with determining whether Allow\_Credit is equal to one 610. If Allow\_Credit is not equal to one, the processing proceeds to the calculation and assignment of the variable values 660. If Allow\_Credit is equal to one, the signal OfferRC is compared to a threshold debt value D 620, where D is equal to DebtCredit if DebtCredit is greater than zero, and D has the value zero if DebtCredit is less than or equal to zero. If OfferRC greater than or equal to the threshold debt value D, processing continues with a comparison of the NeedCredit signal to the value (OfferRC – D) 630; alternatively, if OfferRC is less than D, the decision logic of arithmetic logic 590 determines whether OfferSC is greater or equal to NeedCredit 650. The values assigned to the variables dDebt, dCurrent, and CQCreditReturn 660 depend on the branch taken from the comparison of OfferSC and NeedCredit 650.

**[0050]** Following another branch of the processing of arithmetic logic 590, if NeedCredit is less than or equal to (OfferRC – D) 630, the processing is directed to the

calculation and assignment of the variable values 660 as shown in FIG. 6; otherwise, the processing continues with a comparison of the input signal OfferSC to the value (NeedCredit – (Offer RC – D)) 640. From this comparison of the OfferSC signal 640, processing proceeds to the calculation and assignment of the variable values 660 in accordance with the output branch taken from the comparison as shown in FIG. 6.

[0051] The present invention can be included in an article of manufacture (e.g., one or more computer program products) having, for instance, computer usable media. The media has therein, for instance, computer readable program code means or logic (e.g., instructions, code, commands, etc.) to provide and facilitate the capabilities of the present invention. The article of manufacture can be included as a part of a computer system or sold separately.

[0052] Additionally, at least one program storage device readable by a machine embodying at least one program of instructions executable by the machine to perform the capabilities of the present invention can be provided.

[0053] The flow diagrams depicted herein are just examples. There may be many variations to these diagrams or the steps (or operations) described therein without departing from the spirit of the invention. For instance, the steps may be performed in a differing order, or steps may be added, deleted or modified. All of these variations are considered a part of the claimed invention.

[0054] Although preferred embodiments have been depicted and described in detail herein, it will be apparent to those skilled in the relevant art that various modifications, additions, substitutions and the like can be made without departing from the spirit of the invention and these are therefore considered to be within the scope of the invention as defined in the following claims.